

Nonasymptotic Convergence Rates for Cooperative Learning Over Time-Varying Directed Graphs

Angelia Nedić, Alex Olshevsky and César A. Uribe

Abstract

We study the problem of distributed hypothesis testing with a network of agents where some agents repeatedly gain access to information about the correct hypothesis. The group objective is to globally agree on a joint hypothesis that best describes the observed data at all the nodes. We assume that the agents can interact with their neighbors in an unknown sequence of time-varying directed graphs. Following the pioneering work of Jadbabaie, Molavi, Sandroni, and Tahbaz-Salehi, we propose local learning dynamics which combine Bayesian updates at each node with a local aggregation rule of private agent signals. We show that these learning dynamics drive all agents to the set of hypotheses which best explain the data collected at all nodes as long as the sequence of interconnection graphs is uniformly strongly connected. Our main result establishes a non-asymptotic, explicit, geometric convergence rate for the learning dynamic.

I. INTRODUCTION

Recent years have seen a considerable amount of work on the analysis of distributed algorithms. Nonetheless, the study of distributed decision making and computation can be traced back to the classic papers [1], [2], [3] from the 70s and 80s. Applications of such algorithms range from opinion dynamics analysis, network learning and inference, cooperative robotics, communication networks, to social as well as sensor networks. It is the latter settings of social and sensor networks which is the focus of the current paper.

Interactions among people produce exchange of ideas, opinions, observations and experiences, on which new ideas, opinions, and observations are generated. Analyzing dynamic model of such processes generates insight into human behavior and produce algorithms useful in the sensor networking context.

The authors are with the Coordinated Science Laboratory, University of Illinois, 1308 West Main Street, Urbana, IL 61801, USA, {angelia,aolshev2,cauribe2}@illinois.edu. This research is supported partially by the National Science Foundation under grant no. CCF-1017564 and by the Office of Naval Research under grant no. N00014-12-1-0998.

We consider an agent network where agents repeatedly receive information from their neighbors and private signals from an external source, which provide samples from random variable with unknown distribution. The agents would like to collectively agree on a hypothesis (distribution) that best explains the data.

Initial results on learning in social networks are described in [4], where local update rules are designed such that it matches the Bayes' Theorem. That is, given a prior and new observations, the agent is able to compute likelihood functions in order to generate a new posterior, see [5]. Nevertheless, a fully Bayesian approach might not be possible in general since full knowledge of neither the network structure nor other agents hypothesis might be available [6]. Fortunately, non-Bayesian methods have been shown successful in learning as well. For example, in [7], the authors propose a modification of Bayes' rule that accounts for over-reactions or under-reactions to new information.

In a distributed setting, several groundbreaking papers have described ways agents achieve global behaviors by repeatedly aggregating local information in a network [8], [9], [10]. For example, in distributed hypothesis testing using belief propagation, convergence and dependence of the communication structure were shown [10]. Later, extensions to finite capacity channels, packet losses, delayed communications and tracking were developed [11], [12]. In [9], the authors proved convergence in probability, asymptotic normality of the distributed estimation and provided conditions under which the distributed estimation is as good as a centralized one. Later in [8], almost sure convergence of a non-Bayesian rule based on arithmetic mean was shown for fixed topology graphs. Extensions to information heterogeneity and asymptotic convergence rates have been derived as well [13]. Following [8], other methods to aggregate Bayes estimates in a network have been explored. In [14], geometric means are used for fixed topologies as well, however the consensus and learning steps are separated. The work in [15] extends the results of [8] to time-varying undirected graphs. In [16], local exponential rates of convergence for undirected gossip-like graphs are studied.

In this paper we propose a non-Bayesian learning rule, analyze its consistency and derive a non-asymptotic rate of convergence for time-varying directed graphs. Our first result shows consistency: we show that over time, the protocol learns the hypothesis or set of hypotheses which better explain the data collected by all the nodes. Moreover, our main result provides a geometric, non-asymptotic, and explicit characterization of the rate of convergence which immediately leads to finite-time bounds which scale intelligibly with the number of nodes.

In a simultaneous independent effort, the authors in [17], [18] proposed a similar non-Bayesian learning algorithm where a local Bayes update is followed by a consensus step. In [17], convergence result for

fixed graphs is provided and large deviation convergence rates are given, proving the existence of a random time after which the beliefs will concentrate exponentially fast. In [18], similar probabilistic bounds for the rate of convergence are derived for fixed graphs and comparisons with the centralized version of the learning rule are provided.

This paper is organized as follows. In Section II we describe the model that we study and the proposed update rule. In Section III we analyze the consistency of the information aggregation and estimation models, while in Section IV we establish non-asymptotic convergence rates of the agent beliefs. Some conclusions and future work directions are given in Section VI.

Notation: Upper case letters represent random variables (e.g. X_k), and the corresponding lower case letters for their realizations (e.g. x_k). Subindex will generally indicate the time index. We write as $[A_k]_{ij}$ the i -th row and j -th column entry of matrix A_k . We write A' for the transpose of a matrix A and x' for the transpose of a vector x . We use I for the identity matrix. Bold letters represent vectors which are assumed to be column vectors. The i 'th entry of a vector will be denoted by a superscript i , i.e., $\mathbf{x}_k = [x_k^1, \dots, x_k^n]'$. We write $\mathbf{1}_n$ to denote the all-ones vector of size n . For a sequence of matrices $\{A_t\}$, we let $A_{t_f:t_i} \triangleq A_{t_f} \cdots A_{t_i+1} A_{t_i}$ for all $t_f \geq t_i \geq 0$. We terms "almost surely" and "independent identically distributed" are abbreviated by *a.s.* and *i.i.d.* respectively.

II. PROBLEM SETUP AND MAIN RESULTS

We consider a group of n agents each of which observes a random variable at each time step $k = 1, 2, 3, \dots$. We use S_k^i to denote the random variable whose samples are observed by agent i at time step k . We denote the set of outcomes of the random variable S_k^i by \mathcal{S}^i , and we assume that this set is finite, i.e., $\mathcal{S}^i = \{s_1^i, s_2^i, \dots, s_{m_i}^i\}$ for all $i = 1, \dots, n$. Furthermore, we assume that all S_k^i are i.i.d. and drawn according to some probability distribution $f^i : \mathcal{S}^i \rightarrow [0, 1]$. For convenience, we stack up all S_k^i 's into a single random vector \mathbf{S}_k .

We assume there is a finite set of hypothesis, $\Theta = \{\theta_1, \theta_2, \dots, \theta_m\}$ and there is a probability distribution $l_i(\cdot|\theta)$ for each agent i and hypothesis $\theta \in \Theta$. Intuitively, we think of $l_i(\cdot|\theta)$ as the probability distribution seen by agent i if hypothesis θ were true. Note that, it is not required for the agents to have an hypothesis that is exactly equal to the unknown distribution f^i . The goal of the agents is to agree on an element of Θ that fits all the observations in the network best (in a technical sense to be described soon).

Agents communicate with their neighbors, this communication is modeled as a graph $\mathcal{G}_k = \{V, E_k\}$ composed of a node set $V = \{1, 2, \dots, n\}$ and a set of directed links E_k .

We will refer to probability distributions over Θ as *beliefs* and assume that agent i begins with an

initial belief μ_0^i , which we also refer to as its *prior distribution* or *prior belief*.

This paper focuses in the study of the group dynamics wherein, at time k , each agent i updates its previous belief μ_k^i to a new belief μ_{k+1}^i as follows:

$$\mu_{k+1}^i(\theta) = \frac{\prod_{j=1}^n \mu_k^j(\theta)^{[A_k]_{ij}} l_i(s_{k+1}^i|\theta)}{\sum_{p=1}^m \prod_{j=1}^n \mu_k^j(\theta_p)^{[A_k]_{ij}} l_i(s_{k+1}^i|\theta_p)}, \quad (1)$$

with $[A_k]_{ij} > 0$ when i receives information from j at time k , and else $[A_k]_{ij} = 0$.

The “weight matrices” A_k satisfy some technical connectivity conditions which have been previously used in convergence analysis of distributed averaging and other consensus algorithms [19], [20], [21]. The assumptions on the communication graph are presented next.

Assumption 1 *The graph sequence $\{\mathcal{G}_k\}$ and the matrix sequence $\{A_k\}$ are such that:*

- 1) A_k is row-stochastic with $[A_k]_{ij} > 0$ if $(j, i) \in E_k$.
- 2) A_k has positive diagonal entries, $[A_k]_{ii} > 0$.
- 3) If $[A_k]_{ij} > 0$ then $[A_k]_{ij} > \eta$ for some positive constant η .
- 4) $\{\mathcal{G}_k\}$ is B -strongly connected, i.e., there is an integer $B \geq 1$ such that the graph $\left\{V, \bigcup_{i=kB}^{(k+1)B-1} E_i\right\}$ is strongly connected for all $k \geq 0$.

As a measure for the explanatory quality of the hypotheses in the set Θ we use the Kullback-Leibler divergence between two discrete probability distributions \mathbf{p} and \mathbf{q} :

$$d(\mathbf{p}||\mathbf{q}) = \sum_{i=1}^n p_i \log \left(\frac{p_i}{q_i} \right).$$

Concretely, the quality of hypothesis θ_j for agent i is measured by the Kullback-Leibler divergence $d(f^i(\cdot) || l_i(\cdot|\theta_j))$ between the true distribution of the signals S_k^i and the probability distribution $l_i(\cdot|\theta_j)$ as seen by agent i if hypothesis θ_j were correct. We use the following assumption on the agents’ best hypotheses.

Assumption 2 *The set Θ^* defined as $\Theta^* \triangleq \bigcap_{i=1}^n \Theta_i$, where $\Theta_i = \arg \min_{\theta \in \Theta} d(f^i(\cdot) || l_i(\cdot|\theta))$ for each i , is non-empty.*

Assumption 2 is satisfied if there is some “true state of the world” $\hat{\theta} \in \Theta$ such that each agent i sees distributions generated according to $\hat{\theta}$, i.e., $f^i(\cdot) = l_i(\cdot|\hat{\theta})$. However, this need not be the case for Assumption 2 to hold. Indeed, the assumption is considerably weaker as it merely requires that the set of hypotheses, which provide the “best fits” for each agent, have at least a single element in common.

We will further require the following assumptions on the initial distribution and the likelihood functions. The first of these is sometimes referred to as the Zero Probability Property [22].

Assumption 3 For all agents $i = 1, \dots, n$,

- 1) The prior beliefs on all $\theta^* \in \Theta^*$ are positive, i.e. $\mu_0^i(\theta^*) > 0$ for all $\theta^* \in \Theta^*$.
- 2) There exists an $\alpha > 0$ such that $l_i(s^i|\theta) > \alpha$ for all $s^i \in \mathcal{S}^i$ and $\theta \in \Theta$.

Assumption 3.1 can be relaxed to a requirement that all prior beliefs are positive for some $\theta^* \in \Theta^*$. Both of these conditions are equally complex to be satisfied. They can be satisfied by letting each agent have a uniform prior belief, which is reasonable in the absence of any initial information about the goodness of the hypotheses.

We now state our first result, which asserts that the dynamics in Eq. (1) concentrate all agent's beliefs in the optimal hypothesis set. We provide its proof in Section III.

Theorem 1 Under Assumptions 1, 2, and 3, the update rule of Eq. (1) has the following property:

$$\lim_{k \rightarrow \infty} \mu_k^i(\theta) = 0 \quad a.s. \quad \forall \theta \notin \Theta^*, \quad i = 1, \dots, n.$$

The result states that the agents' beliefs will concentrate on the set Θ^* asymptotically as $k \rightarrow \infty$.

Our main result is a non-asymptotic explicit convergence rate, given in the following theorem, proven in Section IV.

Theorem 2 Let Assumptions 1, 2, and 3 hold. Also, let $\rho \in (0, 1)$ be a given error percentile (or confidence value). Then, the update rule of Eq. (1) has the following property: there exists an integer $N(\rho)$ such that, with probability $1 - \rho$, for all $k \geq N(\rho)$ there holds that for any $\theta \notin \Theta^*$,

$$\begin{aligned} \mu_k^i(\theta) &\leq \exp\left(-\frac{k}{2}\gamma_2 + \gamma_1\right) \quad \forall i = 1, \dots, n, \\ \text{where } N(\rho) &\triangleq \frac{8(\log(\alpha))^2 \log\left(\frac{1}{\rho}\right)}{\gamma_2^2} + 1, \\ \gamma_1 &\triangleq \max_{\substack{\theta^* \in \Theta^* \\ \theta \notin \Theta^*}} \left\{ \max_{1 \leq i \leq n} \log \frac{\mu_0^i(\theta)}{\mu_0^i(\theta^*)} + \frac{C}{1-\lambda} \|\mathbf{H}(\theta)\|_1 \right\}, \\ \gamma_2 &\triangleq \frac{\delta}{n} \min_{\theta \notin \Theta^*} \|\mathbf{H}(\theta)\|_1 \\ [\mathbf{H}(\theta)]_i &= d(f^i(\cdot) \| l_i(\cdot | \theta)) - d(f^i(\cdot) \| l_i(\cdot | \theta^*)), \end{aligned}$$

with α from Assumption 3.2.

The constants C , δ and λ satisfy the following relations:

(1) For general B -connected graph sequences $\{\mathcal{G}_k\}$,

$$C = 2, \quad \lambda \leq (1 - \eta^{nB})^{\frac{1}{B}}, \quad \delta \geq \frac{1}{\eta^{nB}}.$$

(2) If every matrix A_k is doubly stochastic,

$$C = \sqrt{2}, \quad \lambda = \left(1 - \frac{\eta}{4n^2}\right)^{\frac{1}{B}}, \quad \delta = 1.$$

(3) If each \mathcal{G}_k is an undirected graph and each A_k is the lazy Metropolis matrix, i.e. the stochastic matrix which satisfies

$$[A_k]_{ij} = \frac{1}{2 \max(d(i), d(j))} \quad \text{for all } \{i, j\} \in \mathcal{G}_k,$$

then
$$C = \sqrt{2}, \quad \lambda = 1 - \frac{1}{\mathcal{O}(n^2)}, \quad \delta = 1.$$

Note that $\mathbf{H}(\theta)$ does not depend on θ^* since $d(f^i(\cdot) \| l_i(\cdot | \theta^*))$ is the same for all θ^* .

In contrast to the previous literature, this convergence rate is not only geometric but also non-asymptotic and explicit in the sense of immediately leading to bounds which scale intelligible in terms of the number of nodes. For example, in the case of doubly stochastic matrices, Theorem 2 immediately implies that, after a transient time, which scales cubically in the number n of nodes, the network will achieve exponential decay to the correct answer with the exponent $-\frac{1}{2} \min_{\theta^* \in \Theta^*} \|\mathbf{H}(\theta)\|_1 / n$.

Now, consider the case when Assumption 3.1 is relaxed to the following requirement: *The prior beliefs on some $\theta^* \in \Theta^*$ are positive* (i.e. $\mu_0^i(\theta^*) > 0$ for some $\theta^* \in \Theta^*$ and all i). Then, it can be seen that the Theorem 2 is valid with $\max_{\theta^* \in \Theta^*}$ and $\min_{\theta^* \in \Theta^*}$ replaced, respectively, by $\max_{\theta^* \in \tilde{\Theta}^*}$ and $\max_{\theta^* \in \tilde{\Theta}^*}$, where $\tilde{\Theta}^* \subseteq \Theta^*$ is the set of all $\theta^* \in \Theta^*$ for which all the agents priors μ_0^i are positive.

III. CONSISTENCY OF THE LEARNING RULE

In this section we prove Theorem 1, which provides a statement about the consistency (see [23], [24]) of the distributed estimator given in Eq. (1). Our analysis will require some auxiliary results. First, we will recall some results from [25] about the convergence of a product of row stochastic matrices.

Lemma 1 [25], [26] *Under Assumption 1, for a graph sequence $\{\mathcal{G}_k\}$ and each $t \geq 0$, there is a stochastic vector ϕ_t (meaning its entries are nonnegative and sum to one) such that for all i, j and $k \geq t$,*

$$\left| [A_{k:t}]_{ij} - \phi_t^j \right| \leq C \lambda^{k-t} \quad \forall k \geq t \geq 0$$

where $C > 0$ and $\lambda \in (0, 1)$ satisfy the relations described in Theorem 2.

The proof of Lemma 1 may be found in [25], with the exception of the bounds on C, λ for the lazy Metropolis chains which we omit here due to space constraints.

Lemma 2 [25] *Let the graph sequence $\{\mathcal{G}_k\}$ satisfy Assumption 1. Define*

$$\delta \triangleq \inf_{k \geq 0} \left(\min_{1 \leq i \leq n} [\mathbf{1}'_n A_{k:0}]_i \right). \quad (2)$$

Then, $\delta \geq \eta^{nB}$, and if all A_k are doubly stochastic, then $\delta = 1$. Furthermore, the sequence ϕ_t from Lemma 1 satisfies $\phi_t^j \geq \delta/n$ for all $t \geq 0, j = 1, \dots, n$.

Next, we need a technical lemma regarding the weighted average of random variables with a finite variance.

Lemma 3 *If assumptions 1, 2 and 3 hold. Then for a graph sequence $\{\mathcal{G}_k\}$ we have for any $\theta \notin \Theta^*$ and $\theta^* \in \Theta^*$,*

$$\lim_{k \rightarrow \infty} \frac{1}{k} \sum_{t=1}^k A_{k:t} \mathcal{L}_t^\theta + \frac{1}{k} \sum_{t=1}^k \mathbf{1}_n \phi'_t \mathbf{H}(\theta) = 0 \quad a.s.$$

where \mathcal{L}_t^θ is the random vector with coordinates given by

$$[\mathcal{L}_t^\theta]_i = \log \frac{l_i(S_t^i | \theta)}{l_i(S_t^i | \theta^*)} \quad \forall i = 1, \dots, n,$$

while the vector $\mathbf{H}(\theta)$ has coordinates given by $[\mathbf{H}(\theta)]_i = d(f^i(\cdot) || l_i(\cdot | \theta)) - d(f^i(\cdot) || l_i(\cdot | \theta^*))$.

Proof: Adding and subtracting $\frac{1}{k} \sum_{t=1}^k \mathbf{1}_n \phi'_t \mathcal{L}_t^\theta$ yields

$$\frac{1}{k} \sum_{t=1}^k \left(A_{k:t} \mathcal{L}_t^\theta + \mathbf{1}_n \phi'_t \mathbf{H}(\theta) \right) = \frac{1}{k} \sum_{t=1}^k (A_{k:t} - \mathbf{1}_n \phi'_t) \mathcal{L}_t^\theta + \frac{1}{k} \sum_{t=1}^k \mathbf{1}_n \phi'_t (\mathcal{L}_t^\theta + \mathbf{H}(\theta)). \quad (3)$$

By Lemma 1, $\lim_{k \rightarrow \infty} A_{k:t} = \mathbf{1}_n \phi'_t$ for all $t \geq 0$. Moreover, each of the entries of \mathcal{L}_t^θ are upper bounded by Assumption 2. Thus, the first term on the right hand side of Eq. (3) goes to zero as we take the limit over $k \rightarrow \infty$. Regarding the second term in Eq. (3), by the definition of the KL divergence measure, we have that

$$\begin{aligned} \mathbb{E} \left[\log \frac{l_i(S_t^i | \theta)}{l_i(S_t^i | \theta^*)} \right] &= \sum_{j=1}^{m_i} f^i(s_j^i) \log \frac{l_i(s_j^i | \theta)}{l_i(s_j^i | \theta^*)} \\ &= \sum_{j=1}^{m_i} f^i(s_j^i) \log \left(\frac{l_i(s_j^i | \theta)}{l_i(s_j^i | \theta^*)} \frac{f^i(s_j^i)}{f^i(s_j^i)} \right) \end{aligned}$$

$$= d(f^i(\cdot) \| l_i(\cdot | \theta^*)) - d(f^i(\cdot) \| l_i(\cdot | \theta))$$

or equivalently $\mathbb{E}[\mathcal{L}_t^\theta] = -\mathbf{H}(\theta)$.

Kolmogorov's strong law of large numbers states that if $\{X_t\}$ is a sequence of independent random variables with variances such that $\sum_{k=1}^{\infty} \frac{\text{Var}(X_k)}{k^2} < \infty$, then $\frac{1}{n} \sum_{k=1}^n X_k - \frac{1}{n} \sum_{k=1}^n \mathbb{E}[X_k] \rightarrow 0$ a.s. Let $X_t = \phi_t' \mathcal{L}_t^\theta$. Then, by using Assumptions 1 and 3.2, it can be seen that $\sup_{t \geq 0} \text{Var}(X_t) < \infty$.

The result follows by Lemma 1 and Kolmogorov's strong law of large numbers. \blacksquare

With Lemma 3 in place, we are ready to prove Theorem 1. The proof of Theorem 1 (and also Theorem 2) makes use of the following quantities: for all $i = 1, \dots, n$ and $k \geq 0$,

$$\varphi_k^i(\theta) \triangleq \log \frac{\mu_k^i(\theta)}{\mu_k^i(\theta^*)} \quad \text{for all } \theta \in \Theta, \quad (4)$$

defined for any $\theta^* \in \Theta^*$ (dependence on θ^* is suppressed). *Proof: (Theorem 1)* Dividing both sides of (1) by $\mu_{k+1}^i(\theta^*)$, then using the log function and the definition of $\varphi_k^i(\theta)$ we obtain:

$$\varphi_{k+1}^i(\theta) = \sum_{j=1}^n [A_k]_{ij} \varphi_k^j(\theta) + \log \frac{l_i(s_{k+1}^i | \theta)}{l_i(s_{k+1}^i | \theta^*)}.$$

Stacking up the values $\varphi_{k+1}^i(\theta)$ over agents $i = 1, \dots, n$, into a single vector $\boldsymbol{\varphi}_{k+1}(\theta)$, we can compactly write the preceding relations, as follows:

$$\boldsymbol{\varphi}_{k+1}(\theta) = A_k \boldsymbol{\varphi}_k(\theta) + \mathcal{L}_{k+1}^\theta, \quad (5)$$

which implies that for all $k \geq 0$,

$$\boldsymbol{\varphi}_{k+1}(\theta) = A_{k:0} \boldsymbol{\varphi}_0(\theta) + \sum_{t=1}^k A_{k:t} \mathcal{L}_t^\theta + \mathcal{L}_{k+1}^\theta. \quad (6)$$

We add and subtract $\sum_{t=1}^k \mathbf{1}_n \phi_t' \mathbf{H}(\theta)$ in Eq. (6), then

$$\boldsymbol{\varphi}_{k+1}(\theta) = A_{k:0} \boldsymbol{\varphi}_0(\theta) + \sum_{t=1}^k \left(A_{k:t} \mathcal{L}_t^\theta + \mathbf{1}_n \phi_t' \mathbf{H}(\theta) \right) + \mathcal{L}_{k+1}^\theta - \sum_{t=1}^k \mathbf{1}_n \phi_t' \mathbf{H}(\theta).$$

By using the lower bounds on ϕ_t described in Lemma 2 and the fact that $\mathbf{H}(\theta, \theta^*) \geq 0$, we obtain

$$\boldsymbol{\varphi}_{k+1}(\theta) \leq A_{k:0} \boldsymbol{\varphi}_0(\theta) + \sum_{t=1}^k \left(A_{k:t} \mathcal{L}_t^\theta + \mathbf{1}_n \phi_t' \mathbf{H}(\theta) \right) + \mathcal{L}_{k+1}^\theta - \frac{\delta}{n} k \|\mathbf{H}(\theta)\|_1 \mathbf{1}_n.$$

Therefore, we have

$$\lim_{k \rightarrow \infty} \frac{1}{k} \boldsymbol{\varphi}_{k+1}(\theta) \leq \lim_{k \rightarrow \infty} \frac{1}{k} A_{k:0} \boldsymbol{\varphi}_0(\theta) - \frac{\delta}{n} \|\mathbf{H}(\theta)\|_1 \mathbf{1}_n + \lim_{k \rightarrow \infty} \frac{1}{k} \mathcal{L}_{k+1}^\theta + \lim_{k \rightarrow \infty} \frac{1}{k} \sum_{t=1}^k \left(A_{k:t} \mathcal{L}_t^\theta + \mathbf{1}_n \phi_t' \mathbf{H}(\theta) \right).$$

The first term of the right hand side of the preceding relation converges to zero deterministically. The third term goes to zero as well since \mathcal{L}_t^θ is bounded, and the fourth term converges to zero almost surely by Lemma 3. Consequently,

$$\lim_{k \rightarrow \infty} \frac{1}{k} \varphi_{k+1}(\theta) \leq -\frac{\delta}{n} \|\mathbf{H}(\theta)\|_1 \mathbf{1}_n \quad a.s. \quad (7)$$

Now if $\theta \notin \Theta^*$, then $\mathbf{H}(\theta, \theta^*) > 0$ and, thus, $\varphi_k(\theta) \rightarrow -\infty$ almost surely. This implies $\mu_k(\theta) \rightarrow 0$ almost surely. \blacksquare

IV. NON-ASYMPTOTIC RATE OF CONVERGENCE

In this section, we prove Theorem 2, which states an explicit rate of convergence for cooperative agent learning process. Before proving the theorem, we will estate an auxiliary lemma that provides a bound for the expectation of the random variables $\varphi_k^i(\theta)$ as defined in Eq. (4).

Lemma 4 *Let $\theta^* \in \Theta^*$ be arbitrary, and consider $\varphi_k^i(\theta)$ as defined in Eq. (4). Then, for any $\theta \notin \Theta^*$ we have*

$$\mathbb{E}[\varphi_{k+1}^i(\theta)] \leq \gamma_1 - (k+1)\gamma_2 \quad \text{for all } i \text{ and } k \geq 0,$$

where γ_1 and γ_2 are defined in Theorem 2.

Proof: The expected value of Eq. (5) and $\mathbb{E}[\mathcal{L}_{k+1}^\theta] = -\mathbf{H}(\theta)$, gives

$$\mathbb{E}[\varphi_{k+1}(\theta)] = A_k \mathbb{E}[\varphi_k(\theta)] - \mathbf{H}(\theta)$$

Therefore, by recursion we can see that for all $k \geq 0$,

$$\mathbb{E}[\varphi_{k+1}(\theta)] = A_{k:0} \varphi_0(\theta) - \sum_{t=1}^k A_{k:t} \mathbf{H}(\theta) - \mathbf{H}(\theta).$$

By adding and subtracting $\sum_{t=1}^k \mathbf{1}_n \phi_t' \mathbf{H}(\theta)$, we obtain

$$\mathbb{E}[\varphi_{k+1}(\theta)] = A_{k:0} \varphi_0(\theta) + \sum_{t=1}^k (\mathbf{1}_n \phi_t' - A_{k:t}) \mathbf{H}(\theta) - \sum_{t=1}^k \mathbf{1}_n \phi_t' \mathbf{H}(\theta) - \mathbf{H}(\theta).$$

We removed the last term of the right hand side in the preceding relation since $\mathbf{H}(\theta) \geq 0$. Moreover, bounding the entries for the first two terms on the right hand side and using the fact that $A_{k:0}$ is a stochastic matrix, we have that

$$\mathbb{E}[\varphi_{k+1}(\theta)] \leq \|\varphi_0(\theta)\|_\infty \mathbf{1}_n - \sum_{t=1}^k \mathbf{1}_n \phi_t' \mathbf{H}(\theta) + \sum_{t=1}^k \max_{1 \leq i, j \leq n} |\phi_t^j - [A_{k:t}]_{ij}| \|\mathbf{H}(\theta)\|_1 \mathbf{1}_n$$

Next, we use the upper bound on terms $|\phi_t^j - [A_{k:t}]_{ij}|$ from Lemma 1 and the lower bound for the entries in ϕ_t as given in Lemma 2, and we arrive at the following relation:

$$\mathbb{E} [\varphi_{k+1}(\theta)] \leq \|\varphi_0(\theta)\|_\infty \mathbf{1}_n + \sum_{t=1}^k C \lambda^{k-t} \|\mathbf{H}(\theta)\|_1 \mathbf{1}_n - k \frac{\delta}{n} \|\mathbf{H}(\theta)\|_1 \mathbf{1}_n$$

and the result follows. ■

The proof of Theorem 2 uses the McDiarmid's inequality [27]. This will provide bounds on the probability that the beliefs exceed a given value ϵ . McDiarmid's inequality is provided below.

Theorem 3 (McDiarmid's inequality [27]) *Let $\{X_t\}_{t=1}^k = (X_1, \dots, X_k)$ be a sequence of independent random variables with $X_t \in \mathcal{X}$. If a function $g : \{X_t\}_{t=1}^k \rightarrow \mathbb{R}$ has bounded differences, i.e., for all t ,*

$$\sup_{X_t \in \mathcal{X}} g(\dots, X_t, \dots) - \inf_{Y_t \in \mathcal{X}} g(\dots, Y_t, \dots) \leq c_t$$

then for any $\epsilon > 0$ and all $k \geq 1$,

$$\mathbb{P} \left(g(\{X_t\}_{t=1}^k) - \mathbb{E} [g(\{X_t\}_{t=1}^k)] \geq \epsilon \right) \leq \exp \left(\frac{-2\epsilon^2}{\sum_{t=1}^k c_t^2} \right)$$

Now, we are ready to prove Theorem 2.

Proof: (Theorem 2) First we will express the belief $\mu_{k+1}^i(\theta)$ in terms of the variable $\varphi_{k+1}^i(\theta)$. This will allow us to use the McDiarmid's inequality to obtain the concentration bounds. By dynamics of Eq. (1) and Assumption 3.1, since $\mu_{k+1}^i(\theta^*) \in (0, 1]$ for any $\theta^* \in \Theta^*$, we have

$$\mu_{k+1}^i(\theta) \leq \frac{\mu_{k+1}^i(\theta)}{\mu_{k+1}^i(\theta^*)} = \exp(\varphi_{k+1}^i(\theta))$$

Therefore,

$$\begin{aligned} \mathbb{P} \left(\mu_{k+1}^i(\theta) \geq \exp \left(\frac{-k\gamma_2}{2} + \gamma_1 \right) \right) &\leq \mathbb{P} \left(\varphi_{k+1}^i(\theta) \geq \frac{-k\gamma_2}{2} + \gamma_1 \right) \\ &= \mathbb{P} \left(\varphi_{k+1}^i(\theta) - \mathbb{E} [\varphi_{k+1}^i(\theta)] \geq -\frac{k}{2}\gamma_2 + \gamma_1 - \mathbb{E} [\varphi_{k+1}^i(\theta)] \right) \\ &= \mathbb{P} \left(\varphi_{k+1}^i(\theta) - \mathbb{E} [\varphi_{k+1}^i(\theta)] \geq \frac{k}{2}\gamma_2 \right), \end{aligned}$$

where the last equality follows from Lemma 4.

We now view $\varphi_{k+1}^i(\theta)$ a function of the random vectors s_1, \dots, s_k, s_{k+1} , see Eq. (6), where $s_t = (s_t^1, \dots, s_t^n) \in \mathcal{S}$ for all t . Thus, for all t with $1 \leq t \leq k$, we have

$$\max_{s_t \in \mathcal{S}} \varphi_{k+1}^i(\theta) - \min_{s_t \in \mathcal{S}} \varphi_{k+1}^i(\theta) = \max_{s_t \in \mathcal{S}} \sum_{j=1}^n [A_{k:t}]_{ij} [\mathcal{L}_t^\theta]_j - \min_{s_t \in \mathcal{S}} \sum_{j=1}^n [A_{k:t}]_{ij} [\mathcal{L}_t^\theta]_j$$

$$\begin{aligned}
&= \max_{\mathbf{s}_t \in \mathcal{S}} \sum_{j=1}^n [A_{k:t}]_{ij} \log \frac{l_j(s_t^j | \theta)}{l_j(s_t^j | \theta^*)} - \min_{\mathbf{s}_t \in \mathcal{S}} \sum_{j=1}^n [A_{k:t}]_{ij} \log \frac{l_j(s_t^j | \theta)}{l_j(s_t^j | \theta^*)} \\
&\leq \log \frac{1}{\alpha} + \log \frac{1}{\alpha} \\
&= 2 \log \frac{1}{\alpha}.
\end{aligned}$$

Similarly, from Eq. (6) we can see that

$$\max_{\mathbf{s}_{k+1} \in \mathcal{S}} \varphi_{k+1}^i(\theta) - \min_{\mathbf{s}_{k+1} \in \mathcal{S}} \varphi_{k+1}^i(\theta) \leq 2 \log \frac{1}{\alpha}.$$

It follows that $\varphi_{k+1}^i(\theta)$ has bounded variations and by McDiarmid's inequality (Theorem 3) we obtain the following concentration inequality,

$$\begin{aligned}
\mathbb{P} \left(\varphi_{k+1}^i(\theta) - \mathbb{E} [\varphi_{k+1}^i(\theta)] \geq \frac{k}{2} \gamma_2 \right) &\leq \exp \left(\frac{-\frac{1}{2} (k \gamma_2)^2}{\sum_{t=1}^{k+1} (2 \log \frac{1}{\alpha})^2} \right) \\
&= \exp \left(\frac{-(k \gamma_2)^2}{8(k+1) (\log \frac{1}{\alpha})^2} \right) \\
&\leq \exp \left(-\frac{(k-1) \gamma_2^2}{8 (\log \alpha)^2} \right)
\end{aligned}$$

Finally, for a given confidence level ρ , in order to have $\mathbb{P} \left(\mu_k^i(\theta) \geq \exp \left(-\frac{k \gamma_2}{2} + \gamma_1 \right) \right) \leq \rho$ the desired result follows. ■

V. SIMULATION RESULTS

In this section we show simulation results for a group of agents connected over a time-varying directed graph, shown in Figure 1, for some specific weighting matrices. Each agent updates its beliefs according to Eq. (1).

Note that the graph is such that the edge connecting agent 1 and agent 2 is switching on and off at each time step. Agents 2-6 connecting edges are changing at each time step as well.




Fig. 1. Time-Varying graph with a switching external agent

Every agent i receives information from a binary random variable $S_k^i : \Omega \rightarrow \{0, 1\}$ with probability distribution $f^i(0) = 0.1$ and $f^i(1) = 0.9$ for all i 's. Moreover, every agent has two possible models θ_1 and θ_2 . Agent 1 hypotheses have the following likelihood functions: $l_1(0|\theta_1) = 0.2$ and $l_1(1|\theta_1) = 0.8$ for hypothesis θ_1 ; and $l_1(0|\theta_2) = 0.9$ and $l_1(1|\theta_2) = 0.1$ for hypothesis θ_2 . Therefore, hypothesis θ_1 is closer to the true distribution. On the other hand, agents 2 to 6 have uniformly distributed observationally equivalent hypothesis for both θ_1 and θ_2 , that is, they are not able to differentiate between the hypothesis individually. Thus $l_i(s|\theta) = 0.5$ for $i = \{2, \dots, 6\}$, $s = \{0, 1\}$ and $\theta = \{\theta_1, \theta_2\}$.

Figure 2 shows the empirical mean over 5000 Monte Carlo simulations of the beliefs on hypothesis θ_2 of agents 1, 4, 5 and 6. Results show that agent 1 is the fastest learning agent, since is the one with the correct model. Nevertheless, all other agents are converging to the correct parameter model as well, even if they do not have differentiable models.




Fig. 2. Simulation results for Agents 1, 4, 5 and 6

VI. CONCLUSIONS AND FUTURE WORK

We have studied the consistency and the rate of convergence for a distributed non-Bayesian learning system. We have shown almost sure consistency and have provided bounds on the global exponential rate of convergence. The novelty of our results is in the establishment of convergence rate estimates that are non-asymptotic, geometric, and explicit, in the sense that the bounds capture the quantities characterizing the graph sequence properties as well as the agent learning capabilities. This results were derived for general time-varying directed graphs.

Our work suggests a number of open questions. It is natural to attempt to extensions to continuous spaces, on the number of agents, on the number of hypothesis, etc. This result can be extended to tracking problems where the distribution of the observations changes with time. When the number of hypothesis is large, ideas from social sampling can also be incorporated in this framework [28]. Moreover, the possibility of corrupted measurements or conflicting models between the agents are also of interest, especially in the setting of social networks.

REFERENCES

- [1] R. J. Aumann, “Agreeing to disagree,” *The annals of statistics*, pp. 1236–1239, 1976.
- [2] V. Borkar and P. P. Varaiya, “Asymptotic agreement in distributed estimation,” *IEEE Transactions on Automatic Control*, vol. 27, no. 3, pp. 650–655, 1982.
- [3] J. N. Tsitsiklis and M. Athans, “Convergence and asymptotic agreement in distributed decision problems,” *IEEE Transactions on Automatic Control*, vol. 29, no. 1, pp. 42–50, 1984.
- [4] D. Acemoglu, M. A. Dahleh, I. Lobel, and A. Ozdaglar, “Bayesian learning in social networks,” *The Review of Economic Studies*, vol. 78, no. 4, pp. 1201–1236, 2011.
- [5] M. Mueller-Frank, “A general framework for rational learning in social networks,” *Theoretical Economics*, vol. 8, no. 1, pp. 1–40, 2013.
- [6] D. Gale and S. Kariv, “Bayesian learning in social networks,” *Games and Economic Behavior*, vol. 45, no. 2, pp. 329–346, 2003.
- [7] L. G. Epstein, J. Noor, and A. Sandroni, “Non-bayesian learning,” *The BE Journal of Theoretical Economics*, vol. 10, no. 1, 2010.
- [8] A. Jadbabaie, P. Molavi, A. Sandroni, and A. Tahbaz-Salehi, “Non-bayesian social learning,” *Games and Economic Behavior*, vol. 76, no. 1, pp. 210–225, 2012.
- [9] K. Rahnama Rad and A. Tahbaz-Salehi, “Distributed parameter estimation in networks,” in *IEEE Conference on Decision and Control*, 2010, pp. 5050–5055.
- [10] M. Alanyali, S. Venkatesh, O. Savas, and S. Aeron, “Distributed bayesian hypothesis testing in sensor networks,” in *American Control Conference*, vol. 6, 2004, pp. 5369–5374.
- [11] V. Saligrama, M. Alanyali, and O. Savas, “Distributed detection in sensor networks with packet losses and finite capacity links,” *IEEE Transactions on Signal Processing*, vol. 54, no. 11, pp. 4118–4132, 2006.

- [12] R. Rahman, M. Alanyali, and V. Saligrama, "Distributed tracking in multihop sensor networks with communication delays," *IEEE Transactions on Signal Processing*, vol. 55, no. 9, pp. 4656–4668, 2007.
- [13] A. Jadbabaie, P. Molavi, and A. Tahbaz-Salehi, "Information heterogeneity and the speed of learning in social networks," *Columbia Business School Research Paper*, no. 13-28, 2013.
- [14] S. Bandyopadhyay and S.-J. Chung, "Distributed estimation using bayesian consensus filtering," in *American Control Conference (ACC)*, 2014, pp. 634–641.
- [15] Q. Liu, A. Fang, L. Wang, and X. Wang, "Social learning with time-varying weights," *Journal of Systems Science and Complexity*, vol. 27, no. 3, pp. 581–593, 2014.
- [16] S. Shahrampour and A. Jadbabaie, "Exponentially fast parameter estimation in networks using distributed dual averaging," in *IEEE Conference on Decision and Control*, 2013, pp. 6196–6201.
- [17] A. Lalitha, T. Javidi, and A. Sarwate, "Social learning and distributed hypothesis testing," *preprint arXiv:1410.4307*, 2015.
- [18] S. Shahrampour, A. Rakhlin, and A. Jadbabaie, "Distributed detection: Finite-time analysis and impact of network topology," *arXiv preprint arXiv:1409.8606*, 2014.
- [19] D. P. Bertsekas and J. N. Tsitsiklis, *Parallel and distributed computation: numerical methods*. Prentice-Hall, Inc., 1989.
- [20] L. Moreau, "Stability of multiagent systems with time-dependent communication links," *IEEE Transactions on Automatic Control*, vol. 50, no. 2, pp. 169–182, 2005.
- [21] A. Jadbabaie, J. Lin, and A. S. Morse, "Coordination of groups of mobile autonomous agents using nearest neighbor rules," *IEEE Transactions on Automatic Control*, vol. 48, no. 6, pp. 988–1001, 2003.
- [22] C. Genest, J. V. Zidek *et al.*, "Combining probability distributions: A critique and an annotated bibliography," *Statistical Science*, vol. 1, no. 1, pp. 114–135, 1986.
- [23] J. L. Doob, "Application of the theory of martingales," *Le calcul des probabilités et ses applications*, pp. 23–27, 1949.
- [24] S. Ghosal, "A review of consistency and convergence of posterior distribution," in *Varanashi Symposium in Bayesian Inference*, Banaras Hindu University, 1997.
- [25] A. Nedic and A. Olshevsky, "Distributed optimization over time-varying directed graphs," *IEEE Transactions on Automatic Control*, vol. 60, no. 3, pp. 601–615, 2015.
- [26] A. Nedić, A. Olshevsky, A. Ozdaglar, and J. N. Tsitsiklis, "On distributed averaging algorithms and quantization effects," *IEEE Transactions on Automatic Control*, vol. 54, no. 11, pp. 2506–2517, 2009.
- [27] C. McDiarmid, "On the method of bounded differences," *Surveys in combinatorics*, vol. 141, no. 1, pp. 148–188, 1989.
- [28] A. Sarwate and T. Javidi, "Distributed learning of distributions via social sampling," *IEEE Transactions on Automatic Control*, vol. 60, no. 1, pp. 34–45, 2015.